

Signed Distance Function Representation, Tracking, and Mapping

Tanner Schmidt

Overview


- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- SDF limitations
- Related research
 - KinectFusion
 - KinTinuus
 - BundleFusion
 - DART
 - DynamicFusion

Overview

- **Explicit and implicit surface representations**
- SDF fusion
- SDF tracking
- SDF limitations
- Related research
 - KinectFusion
 - KinTinuus
 - BundleFusion
 - DART
 - DynamicFusion

Explicit Surface Representations

- Geometry is stored **explicitly** as a list of points, triangles, or other geometric fragments
 - e.g. meshes, point clouds




Vertices: [(x_0, y_0, z_0) , (x_1, y_1, z_1) , ..., (x_n, y_n, z_n)]

Indices: [(i_0, i_1) , (i_2, i_3) , ..., (i_{n-1}, i_n)]

Implicit Surface Representation


- Geometry is not stored explicitly but rather defined as a level set of a function defined over the space in which the geometry is embedded
 - There are **parametric** representations:



$$f(x, y) = x^2 + y^2 - r^2$$


Implicit Surface Representation

- Geometry is not stored explicitly but rather defined as a level set of a function defined over the space in which the geometry is embedded
 - And there are **nonparametric** representations:



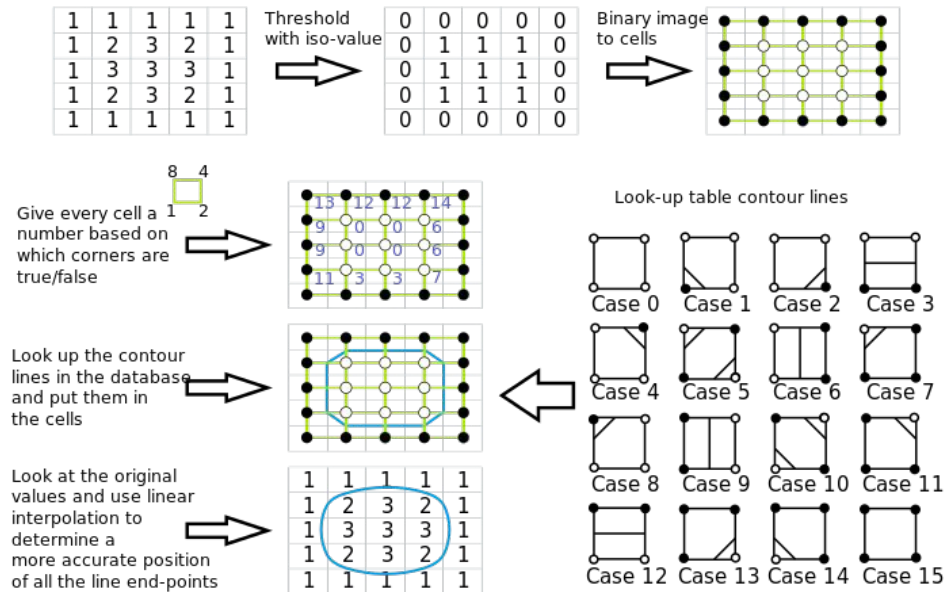
Implicit Surface Representation

- Geometry is not stored explicitly but rather defined as a level set of a function defined over the space in which the geometry is embedded
 - And there are **nonparametric** representations:




Implicit to Explicit Conversion

- In two dimensions, we can use an algorithm called **marching squares**




Implicit to Explicit Conversion in 3D

- Typically done using **marching cubes**, a 3D analogue to marching squares




Implicit to Explicit Conversion in 3D

- Can also be done by **raycasting** for a view-dependent partial surface



Explicit to Implicit Conversion


- Can be done by finding the closest point between each discrete location and any part of the geometry





Explicit to Implicit Conversion

- Can also be done with a **distance transform**

What if we want to build surface representations from raw (noisy) observations?



$$f(x, y)$$



$$\mathcal{D}(f(x, y)) = \min_{x', y'} f(x', y') + (x - x')^2 + (y - y')^2$$

$$\mathcal{D}(f(x, y)) = \min_{x'} (x - x') + \min_{y'} f(x', y') + (y - y')^2$$

Overview

- Explicit and implicit surface representations
- **SDF fusion**
- SDF tracking
- SDF limitations
- Related research
 - KinectFusion
 - KinTinuuous
 - BundleFusion
 - DART
 - DynamicFusion

Signed Distance Function Fusion


$$F_{\mathbf{R}_k}(\mathbf{p}) = \Psi \left(\lambda^{-1} \|(\mathbf{t}_{g,k} - \mathbf{p}\|_2 - \mathbf{R}_k(\mathbf{x}) \right) ,$$

$$\lambda = \|\mathbf{K}^{-1} \dot{\mathbf{x}}\|_2 ,$$


$$\mathbf{x} = \left\lfloor \pi \left(\mathbf{K} \mathbf{T}_{g,k}^{-1} \mathbf{p} \right) \right\rfloor ,$$

$$\Psi(\eta) = \begin{cases} \min \left(1, \frac{\eta}{\mu} \right) \text{sgn}(\eta) & \text{iff } \eta \geq -\mu \\ \text{null} & \text{otherwise} \end{cases}$$

Signed Distance Function Fusion




Signed Distance Function Fusion




1.356	1.384	1.422	1.511	1.607
-------	-------	-------	-------	-------

Signed Distance Function Fusion




1.356	1.384	1.422	1.511	1.607
-------	-------	-------	-------	-------


Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion




Signed Distance Function Fusion



Signed Distance Function Fusion




Signed Distance Function Fusion



Signed Distance Function Fusion


- This addition requires the per-frame projected truncated signed distance volumes to be globally registered




Overview

- Explicit and implicit surface representations
- SDF fusion
- **SDF tracking**
- SDF limitations
- Related research
 - KinectFusion
 - KinTinuuous
 - BundleFusion
 - DART
 - DynamicFusion


Signed Distance Function Tracking




Signed Distance Function Tracking




Signed Distance Function Tracking




Signed Distance Function Tracking




Signed Distance Function Tracking




Signed Distance Function Tracking




Signed Distance Function Tracking




Signed Distance Function Tracking




Point-plane Iterative Closest Point (ICP)




Point-plane Iterative Closest Point (ICP)



Point-plane Iterative Closest Point (ICP)




Point-plane Iterative Closest Point (ICP)




$$E(\theta) = \sum_i \frac{1}{2} \left\| \underline{n_i^{\text{pred}}} \cdot (T_{gl}(\theta) \underline{x_i^{\text{obs}}} - \underline{x_i^{\text{pred}}}) \right\|_2^2$$

Point-plane Iterative Closest Point (ICP)



$$\frac{\partial E(\theta)}{\partial \theta} = \sum_i n_i^{\text{pred}} \cdot \frac{\partial T_{gl}(\theta) x_i^{\text{obs}}}{\partial \theta}$$

Point-plane Iterative Closest Point (ICP)




$$\frac{\partial E(\theta \oplus \Delta\theta)}{\partial \Delta\theta} = \sum_i n_i^{\text{pred}} \cdot \frac{\partial}{\partial \Delta\theta} T_{gl}(\theta \oplus \Delta\theta) x_i^{\text{obs}}$$

Point-plane Iterative Closest Point (ICP)


$$\frac{\partial}{\partial \Delta \theta} T(\theta \oplus \Delta \theta) \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & z & -y \\ 0 & 1 & 0 & -z & 0 & x \\ 0 & 0 & 1 & y & -x & 0 \end{bmatrix}$$

$$T(\theta \oplus \Delta \theta) = e^{\Delta \theta} T(\theta)$$


Point-plane Iterative Closest Point (ICP)



Point-plane Iterative Closest Point (ICP)




Direct Signed Distance Function Tracking




$$E(\theta) = \sum_i \frac{1}{2} \|SDF(T_{gl}(\theta)x_i^{\text{obs}})\|_2^2$$

Direct Signed Distance Function Tracking



$$\frac{\partial E(\theta \oplus \Delta\theta)}{\partial \Delta\theta} = \sum_i \nabla SDF(T_{gl^{t-1}}(\theta) x_i^{\text{obs}}) R_{gl^{t-1}} \frac{\partial T_{l^{t-1}l^t}(\Delta\theta)}{\partial \Delta\theta}$$

Direct Signed Distance Function Tracking



Online fusion

- Tracking requires the fused SDF volume for all frames up to the current frame

Online fusion

- Tracking requires the fused SDF volume for all frames up to the current frame
- We must maintain a **running average** SDF value at each cell




Online fusion

- Tracking requires the fused SDF volume for all frames up to the current frame
- We must maintain a **running average** SDF value at each cell
- Each cell stores both an SDF value and a weight

$$SDF^{0:t} \leftarrow \frac{SDF^{0:t-1}w^{0:t-1} + SDF^tw^t}{w^{0:t-1} + w^t}$$

$$w^{0:t} \leftarrow w^{0:t-1} + w^t$$

Truncated Signed Distance Function




$$\Psi(\eta) = \begin{cases} \min\left(1, \frac{\eta}{\mu}\right) \text{sgn}(\eta) & \text{iff } \eta \geq -\mu \\ \text{null} & \text{otherwise} \end{cases}$$

Overview

- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- **SDF limitations**
- Related research
 - KinectFusion
 - KinTinuus
 - BundleFusion
 - DART
 - DynamicFusion


SDF Limitations

- **Size:**
 - In general, scales linearly with the volume of the computed distance field, and exponentially in the resolution
- **Thin surface representation:**
 - If the resolution is insufficient, observations from opposing sides of thin surfaces can “cancel out”



SDF Limitations

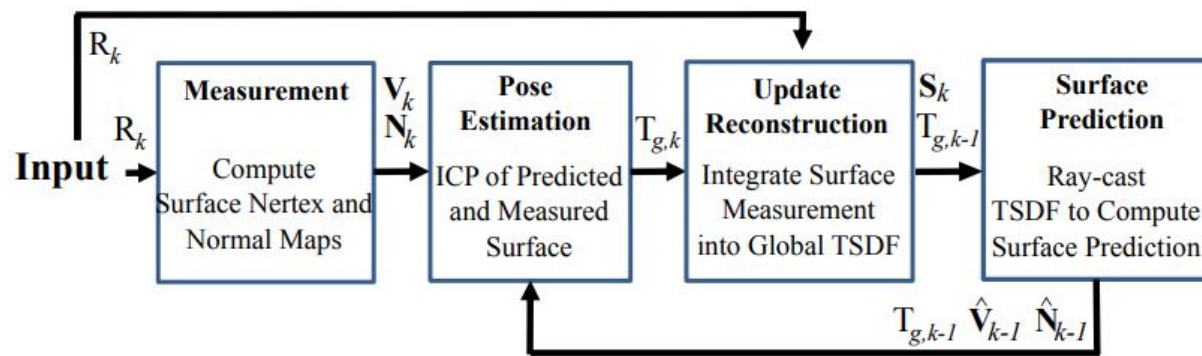
- **Size:**
 - In general, scales linearly with the volume of the computed distance field, and exponentially in the resolution
- **Thin surface representation:**
 - If the resolution is insufficient, observations from opposing sides of thin surfaces can “cancel out”



Overview

- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- SDF limitations
- **Related research**
 - **KinectFusion**
 - KinTinuous
 - BundleFusion
 - DART
 - DynamicFusion

KinectFusion [Newcombe et al.]






Kinect RGB



Raw Normal Map



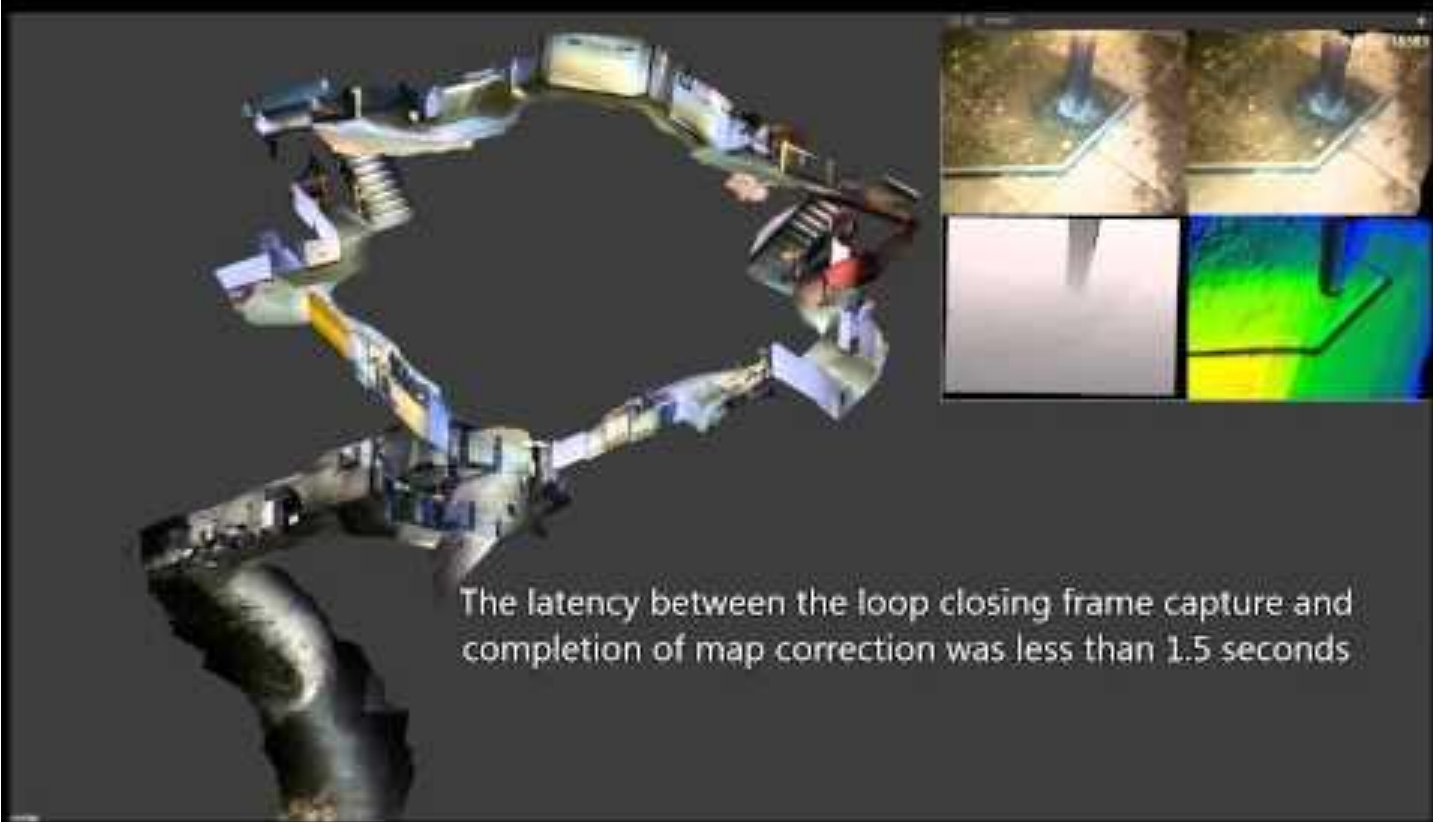
Composited Kinect RGB



**3D Reconstruction +
GPU physics simulation**

Overview

- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- Related research
 - KinectFusion
 - **KinTinuuous**
 - BundleFusion
 - DART
 - DynamicFusion



Overview

- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- Related research
 - KinectFusion
 - KinTinuous
 - **BundleFusion**
 - DART
 - DynamicFusion

Resulting Scan



Overview

- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- Related research
 - KinectFusion
 - KinTinuuous
 - BundleFusion
 - **DART**
 - DynamicFusion



Obs. depth



Obs. to model DA



Obs. to model error



Model depth



Model to obs. DA




Model to obs. error



- The full objective function is given by:

$$\hat{\theta} = \arg \min_{\theta} \sum_{u \in \mathcal{U}} \text{SDF}_{\text{mod}}(\mathbf{x}_u; \theta)^2 + \lambda \sum_{u \in \mathcal{U}} \text{SDF}_{\text{obs}}(\hat{\mathbf{x}}_u(\theta); I)^2$$



Kinect RGB



Tracking



Grasp Planner



Overview

- Explicit and implicit surface representations
- SDF fusion
- SDF tracking
- Related research
 - KinectFusion
 - KinTinuuous
 - BundleFusion
 - DART
 - **DynamicFusion**



Live Input Depth Map



Live Model Output



Live RGB Image (unused)



Canonical Model Reconstruction



Warped Model